arXiv:1512.04499v1 [stat.ME] 14 Dec 2015

False discovery rate control for identifying simultaneous signals

Sihai Dave Zhao

December 15, 2015

Abstract

It is frequently of interest to jointly analyze multiple sequences of multiple tests in order to identify simultaneous signals, defined as features tested in two or more independent studies that are significant in each. For example, researchers often wish to discover genetic variants that are significantly associated with multiple traits. This paper proposes a false discovery rate control procedure for identifying simultaneous signals in two studies. A pair of test statistics is available for each feature, and the goal is to identify features for which both are non-null. Error control is difficult due to the composite nature of a non-discovery, as one of the tests in the pair can still be non-null. Very few existing methods have high power while still provably controlling the false discovery rate. This paper proposes a simple, fast, tuning parameter-free nonparametric procedure that can be shown to provide asymptotically conservative false discovery rate control. Surprisingly, the procedure does not require knowledge of either the null or the alternative distributions of the test statistics. In simulations, the proposed method had higher power and better error control than existing procedures. In an analysis of genome-wide association study results from five psychiatric disorders, it identified more pairs of disorders that share simultaneously significant genetic variants, as well as more variants themselves, compared to other methods. The proposed method is available in the R package `ssa`.

1 Introduction

This paper proposes a new false discovery rate control procedure for identifying simultaneous signals, defined as features, such as genes or single nucleotide polymorphisms in the genomics setting, that are simultaneously important in two separate studies. Specifically, define the signal indicator $I_{kj} \in \{0, 1\}$ to equal 1 if feature j is truly significant in study k and 0 otherwise, $j = 1, \ldots, p$ and $k = 1, 2$. A simultaneous signal is defined as a feature j with $I_{1j} = I_{2j} = 1$. The I_{kj} are unobserved population-level parameters; in practice a pair of test statistics is observed for each feature j, one corresponding to I_{1j} and the other to I_{2j}.

Simultaneous signal identification arises frequently in many different contexts. An important example is in integrative genomics research. In genetics, for instance, it is frequently of interest to identify polymorphisms that are associated with multiple related conditions.

These studies of what is termed genetic pleiotropy are common in recent autoimmune, psychiatric, and cardiovascular disease research (Cross-Disorder Group of the Psychiatric Genomics Consortium et al., 2013a,b; Andreassen et al., 2013b; Li et al., 2015b,a; Rankinen et al., 2015). As another example, identifying findings that replicate across independent studies is a crucial component of reproducible research (Ioannidis et al., 2001; Visscher and Montgomery, 2009; Bogomolov and Heller, 2013; Heller et al., 2014; Heller and Yekutieli, 2014; Donoho, 2015).

Table 1: Results of a simultaneous signal discovery procedure

(I_{1j}, I_{2j})	Not discovered	Discovered	Total	Proportion
(0,0)	U_{00}	V_{00}	p_{00}	π_{00}
(1,0)	U_{10}	V_{10}	p_{10}	π_{10}
(0,1)	U_{01}	V_{01}	p_{01}	π_{01}
(1,1)	U_{11}	V_{11}	p_{11}	π_{11}
Total	$p - R$	R	p	1

This paper develops a method to discover as many simultaneous signals as possible while controlling the rate of false discoveries. Table 1 illustrates the results of applying a discovery procedure to two sequences of test statistics. Define the *false simultaneous discovery rate* as

$$\text{FSDR} = E\left(\frac{V_{00} + V_{10} + V_{01}}{1 \vee R}\right). \tag{1}$$

This paper develops a method to maximize R while maintaining $\text{FSDR} \leq \alpha$.

Controlling FSDR is challenging because a non-simultaneous signal j can still have (I_{1j}, I_{2j}) equal to either $(1, 0)$ or $(0, 1)$. The corresponding test statistics therefore have a composite non-discovery distribution, because their true alternative distributions when $I_{kj} = 1$ are unknown. This makes it difficult to estimate the number of false discoveries incurred by the discovery procedure, a crucial step in false discovery rate control. Very few existing methods are able to have high power to discover simultaneous signals while still provably controlling FSDR; Section 2 gives a review of recent work.

To overcome the composite non-discovery distribution problem, this paper proposes a new nonparametric estimate of the FSDR. One surprising consequence is that the resulting procedure can even be used when the observed test statistics have unknown null distributions; see Section 3.6. This is particularly useful in genomics applications, where test statistics are frequently functions of data with unknown correlation structures. The proposed method is tuning parameter-free, extremely computationally efficient, powerful, and can be shown to provide asymptotically conservative control of FSDR.

2 Related work

The simplest approach to identifying simultaneous signals is to use standard false discovery rate control procedures, e.g., Benjamini and Hochberg (1995), to discover significant features separately in each of the two sequences of test statistics. Features discovered in both

sequences are identified as the simultaneous signals. This ad-hoc procedure is ubiquitous in the scientific literature but cannot control the overall FSDR.

Another common approach is to summarize the pair of statistics for each feature into a single scalar statistic. This reduces the problem to a single sequence of multiple tests, so that a standard false discovery rate control method can be applied. One example is to convert the test statistics into p-values and then take their maximum. However, each summarized statistic has a composite null distribution, which can be difficult to estimate. In addition, it is unclear how to choose the best summary function. Phillips and Ghosh (2014) considered four different possibilities, including taking the maximum p-value. Some were empirically found to be unable to properly control the false discovery rate.

One way around the composite distribution problem is to impose parametric assumptions on the null and alternative distributions of the test statistics. For example, Chung et al. (2014) assumed that p-values were available and were Beta-distributed under the alternative. They used the EM algorithm to estimate the unknown parameters and the direct posterior probability approach of Newton et al. (2004) to control the false discovery rate. Such methods are sensible but the parametric assumptions are restrictive.

A more principled approach is to treat the observed data as a single sequence of bivariate test statistics, as in Chi (2008), Du and Zhang (2014), and Heller and Yekutieli (2014). They all showed that the local false discovery rate (Efron, 2010b) is the optimal scalar summary of the bivariate test statistics. However, this is difficult to calculate in practice. Chi (2008) proposed a Taylor expansion approximation, Du and Zhang (2014) used a single-index model approximation, and Heller and Yekutieli (2014) employed an empirical Bayes approach (Efron, 2010b). These procedures can have high power but tend to be complicated to describe, implement, and analyze theoretically.

3 Proposed procedure

3.1 Statistical setting

Suppose two p-values P_{1j} and P_{2j} are observed for each feature j, $j = 1, \ldots, p$. In a typical genetics example, P_{kj} might be the p-value of the association between the jth genetic variant and the phenotype in the kth genome-wide association study. In practice the two p-values are typically independent because the studies from which they are calculated are conducted on two different groups of subjects. Thus the p-values are modeled as

$$P_{kj} \mid I_{kj} = 0 \sim F^0_{kj} = Unif(0,1), \quad P_{kj} \mid I_{kj} = 1 \sim F^1_{kj}, \quad F^1_{kj} > F^0_{kj}, \quad P_{1j} \perp\!\!\!\perp P_{2j}, \qquad (2)$$

where the I_{kj} are the signal indicators defined in Section 1 and F^0_{kj} and F^1_{kj} are null and alternative distribution functions, respectively. Different features may have different alternative distributions F^1_{kj}. The assumption that the P_{kj} are stochastically smaller under F^1_{kj} is reasonable for two-tailed p-values. For now the p-values will be assumed to be independent across j as well; dependent p-values will be considered in Section 7. Model (2) is a straightforward extension of standard assumptions in the multiple-testing literature (Storey and Tibshirani, 2003; Storey et al., 2004; Donoho and Jin, 2004; Genovese and Wasserman, 2004; Sun and Cai, 2007; Efron, 2010b) to two independent sequences of p-values.

The proposed procedure follows the general three-step false discovery rate control framework outlined by Storey et al. (2004).

3.2 Step 1: define a family of rejection regions

In the standard false discovery rate control problem with a single sequence of p-values, it is natural to use the family of intervals $\{[0, t_1] : t_1 \in [0, 1]\}$ as rejection regions. In the simultaneous discovery problem, however, there are an infinite number of possible shapes that a rejection region can take. In fact, as mentioned in Section 2, it has been shown that the optimal rejection regions are level curves of the local false discovery rate.

The proposed method uses rectangular regions

$$\{[0, t_1] \times [0, t_2] : t_1, t_2 \in [0, 1]\}.$$

The jth feature is identified as a simultaneous signal if $P_{kj} \leq t_k, k = 1, 2$. Du and Zhang (2014) briefly mentioned rectangular regions but did not use them in their final procedure. Chi (2008) studied rectangular regions but limited attention only to rectangles with prespecified aspect ratios. In contrast, the procedure proposed here learns the optimal aspect ratio from the data. While rectangular regions are sub-optimal, they are simple to implement and interpret and have a major advantage over other shapes, which is discussed in Step 2.

3.3 Step 2: estimate the FSDR incurred by a rejection region

This step is where the composite non-discovery distribution causes problems. To see this, assume for now that the F^0_{kj} and F^1_{kj} are all the same across j and equal to some F^0_k and F^1_k, respectively. Then a reasonable estimate of the FSDR incurred by the rectangular rejection region $[0, t_1] \times [0, t_2]$ would be

$$\frac{p\{\pi_{00}F^0_1(t_1)F^0_2(t_2) + \pi_{01}F^0_1(t_1)F^1_2(t_2) + \pi_{10}F^1_1(t_1)F^0_2(t_2)\}}{1 \vee \sum_j I(P_{1j} \leq t_1, P_{2j} \leq t_2)},$$

with $\pi_{ab}, a, b = 0, 1$ defined in Table 1. The numerator is equal to the expected number of non-simultaneous signals falling into the rejection region. This is a straightforward extension of the standard false discovery rate estimator for a single sequence of p-values (Benjamini and Hochberg, 1995; Storey et al., 2004).

The challenge is that the numerator involves F^1_1 and F^1_2, the unknown alternative distributions of the P_{kj}. These quantities do not appear in the classical false discovery rate estimator. Existing methods either impose parametric assumptions, which are restrictive, or attempt to estimate the F^1_k, which can be difficult and time-consuming.

It turns out that this difficulty can be avoided using a simple inequality. Define $F_k(t_k) = (1 - \pi_k)F^0_k(t_k) + \pi_k F^1_k(t_k)$, where $\pi_k = p^{-1}\sum_j I_{kj}$, the proportion of signals in study k. For each k, F_k can be thought of as a marginal distribution function from which the P_{kj} are independently drawn. Some algebra shows that under model (2), when $\pi_{11} > \pi_1\pi_2$,

$$F_1(t_1)F_2(t_2) \geq \{\pi_{00}F^0_1(t_1)F^0_2(t_2) + \pi_{01}F^0_1(t_1)F^1_2(t_2) + \pi_{10}F^1_1(t_1)F^0_2(t_2)\}. \tag{3}$$

This motivates the following proposed estimator for FSDR:

$$\widehat{\text{FSDR}}(t_1, t_2) = \frac{p^{-1} \sum_j I(P_{1j} \leq t_1) \sum_i I(P_{2j} \leq t_2)}{1 \vee \sum_j I(P_{1j} \leq t_1, P_{2j} \leq t_2)}. \tag{4}$$

This nonparametric estimator avoids the composite distribution problem and is tuning parameter-free, making it simple to implement. The advantage of the rectangular rejection regions chosen in Step 1 is that they allow for this simple estimator, as inequality (3) heavily depends on their rectangular shape.

3.4 Step 3: compute the optimal rejection region

The optimal rectangular rejection region discovers the largest number of putative simultaneous signals while maintaining an estimated FSDR less than some desired α:

$$\mathcal{T} = \underset{t_1, t_2 \in \Pi}{\arg\max} \left\{ \hat{G}(t_1, t_2) : \widehat{\text{FSDR}}(t_1, t_2) \leq \alpha \right\}, \tag{5}$$

where $\hat{G}(t_1, t_2) = p^{-1} \sum_j I(P_{1j} \leq t_1, P_{2j} \leq t_2)$ and $\Pi = \{(P_{1j}, P_{2j'}) : 1 \leq j, j' \leq p\}$, the Cartesian product of the two sequences of observed p-values. If there are no thresholds such that $\widehat{\text{FSDR}}(t_1, t_2) \leq \alpha$, then $\mathcal{T}$ can be set to $\{(0, 0)\}$.

In the two-sequence simultaneous discovery problem, an interesting situation arises here as the result of using rectangular rejection regions: in general, (5) does not have a unique solution, because both the objective and the constraint are discrete functions. Thus $\mathcal{T}$ usually contains multiple optimal thresholds, which can identify distinct sets of putative simultaneous signals. Figure 1 illustrates an example where $p = 1,000$, $\pi_{00} = 0.985$, $\pi_{10} = \pi_{01} = \pi_{11} = 0.005$, and the non-null P_{ij} are two-sided p-values corresponding to $N(9, 1)$ test statistics. The figure plots three different optimal rectangular rejection regions at the $\alpha = 0.05$ level, and each region rejects a different set of three features.

This non-uniqueness does not impact the theoretical properties of the proposed method. In fact Theorem 2 in Section 4 shows that any $(\hat{t}_1, \hat{t}_2) \in \mathcal{T}$ can asymptotically control FSDR. On the other hand, in practice it is useful to have a rule for choosing a unique rectangular rejection region. This requires some care. One approach might be to choose the region with the smallest estimated FSDR. However, this does not guarantee uniqueness, as the three regions in Figure 1 all have estimated rates equal to 0.05. Another might be to take the union of the optimal rejection regions. However, the union is no longer necessarily rectangular, so the conservativeness of $\widehat{\text{FSDR}}$ is no longer guaranteed. A third possibility would be to take the intersection, which would be both unique and rectangular. However, this can be very conservative: in the example in Figure 1, the intersection of the three regions contains only one rejected test pair, whereas the individual regions reject three.

Instead, the method proposed here selects the threshold in $\mathcal{T}$ that gives the rectangular region with the largest geometric area:

$$(\hat{t}_1^\star, \hat{t}_2^\star) = \underset{(\hat{t}_1, \hat{t}_2) \in \mathcal{T}}{\arg\max}\, \hat{t}_1 \hat{t}_2. \tag{6}$$

This is intuitively reasonable, avoids the issues mentioned above, and guarantees that $(\hat{t}_1^\star, \hat{t}_2^\star)$ will be unique with probability 1.

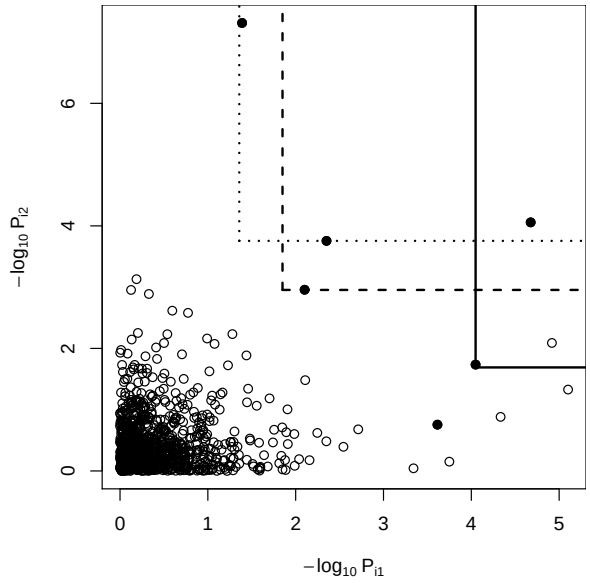

Figure 1: The solid, dashed, and dotted lines demarcate three different estimated rejection regions in $\mathcal{T}$ (5). Filled circles denote true simultaneous associations, with $I_{1j} = I_{2j} = 1$. Each of the three rejection regions has $\widehat{\text{FSDR}} = 0.05$. To make visualization easier, the (P_{1j}, P_{2j}) have been plotted using $-\log_{10}$ p-values.

3.5 Implementation

A simple algorithm for obtaining the thresholds $(\hat{t}_1^\star, \hat{t}_2^\star)$ (6) requires only $O(p^2)$ operations: the P_{1j} and P_{2j} are first sorted using quicksort, which on average requires $O(p \log p)$ operations and at most requires $O(p^2)$. Letting $P_{1(j)}$ and $P_{2(j)}$ denote the resulting order statistics, the algorithm proceeds as follows:

1. Initialize $\hat{t}_1$ and $\hat{t}_2$ to zero. The algorithm considers possible thresholds $(P_{1(k)}, P_{2(l)})$ by looping over $l = 1, \ldots, p$ for each $k = 1, \ldots, p$.

2. Let l' be such that $P_{2l'} = P_{2(l)}$. Then the bivariate empirical distribution function $\hat{G}_{kl} = \hat{G}(P_{1(k)}, P_{2(l)})$ is equal to $\hat{G}_{k,l-1} + p^{-1} I(P_{1l'} \leq P_{1(k)})$. Therefore the estimated false simultaneous discovery rate $\widehat{\text{FSDR}}(P_{1(k)}, P_{2(l)}) = p^{-1} kl / (1 \vee p\hat{G}_{kl})$.

3. Let $\hat{G}(\hat{t}_1, \hat{t}_2)$ be the largest $\hat{G}$ so far. For k and l such that $\widehat{\text{FSDR}}(P_{1(k)}, P_{2(l)}) \leq \alpha$,

   (a) if $\hat{G}_{kl} > \hat{G}(\hat{t}_1, \hat{t}_2)$, update $\hat{t}_1 = P_{1(k)}$ and $\hat{t}_2 = P_{2(l)}$ and store the product $\hat{t}_1 \hat{t}_2$;

   (b) if $\hat{G}_{kl} = \hat{G}(\hat{t}_1, \hat{t}_2)$, update $\hat{t}_1$ and $\hat{t}_2$ only if $P_{1(k)} P_{2(l)}$ is larger than the stored value of $\hat{t}_1 \hat{t}_2$.

To further reduce the computation time, the algorithm does not need to loop over all p^2 elements of Π. Since under model (2) the p-values are stochastically smaller under the alternatives, the optimal $\hat{t}_1^\star$ and $\hat{t}_2^\star$ are likely to be found among the small P_{1i} and P_{2i}. Therefore it suffices to loop over $k = 1, \ldots, m_1$ and $l = 1, \ldots, m_2$, where $m_1, m_2 < p$. As

an example, suppose $p = 10^6$. It is likely that $\hat{t}_1^\star$ and $\hat{t}_2^\star$ lie within the smallest 10% of the P_{1i} and P_{2i}. With $m_1 = m_2 = 10^5$, the algorithm can calculate the optimal thresholds in 115 seconds on a desktop with a 3.2 GHz Intel Core i5 processor with 8 GB RAM. This algorithm has been implemented in the R package `ssa`.

3.6 Discovery using test statistics with unknown null distributions

Though the proposed method was introduced in Section 3.1 using p-values, an examination of the three steps of the procedure reveals that it never uses the fact that the null distribution F_k^0 of the P_{kj} is uniform. In the single-sequence problem, knowledge of the null distribution is necessary to estimate the false discovery rate. In contrast, here in the two-sequence problem the estimator $\widehat{\mathrm{FSDR}}$ is entirely nonparametric, as is the objective function $\hat{G}$ of the threshold optimization problem (5).

This suggests that the proposed procedure can in fact control FSDR for any continuous test statistics T_{1j} and T_{2j} satisfying

$$T_{kj} \mid I_{kj} = 0 \sim F_{kj}^0, \quad T_{kj} \mid I_{kj} = 1 \sim F_{kj}^1, \quad F_{kj}^1 < F_{kj}^0, \quad T_{1j} \perp\!\!\!\perp T_{2j}, \tag{7}$$

where both the alternatives F_{kj}^1 and the nulls F_{kj}^0 are unknown, so that p-values cannot be calculated. This is particularly useful in genomics applications, where sets of correlated genomic features, such as genetic polymorphisms or genes, are often grouped together for analysis. The correlation structure between the features is unknown, resulting in test statistics with unknown nulls. This can also arise in studies with small samples sizes, where asymptotic null distributions may not be very accurate.

The stochastic ordering assumption in model (7) is the opposite of that in model (2): here, the T_{kj} are assumed to be larger under the alternative than under the null, which is in accord with the behavior of standard two-tailed test statistics such as the χ^2 statistic. This requires some minor modifications to the proposed procedure, mostly involving reversing some of the inequalities. The estimator (4) now becomes

$$\widehat{\mathrm{FSDR}}(t_1, t_2) = \frac{p^{-1} \sum_j I(P_{1j} \geq t_1) \sum_i I(P_{2j} \geq t_2)}{1 \vee \sum_j I(P_{1j} \geq t_1, P_{2j} \geq t_2)},$$

the optimal thresholds (5) are now calculated according to

$$\mathcal{T} = \operatorname*{arg\,max}_{t_1, t_2 \in \Pi} \left\{ p^{-1} \sum_j I(P_{1j} \geq t_1, P_{2j} \geq t_2) : \widehat{\mathrm{FSDR}}(t_1, t_2) \leq \alpha \right\},$$

and the uniquely optimal thresholds (6) are now defined as $(\hat{t}_1^\star, \hat{t}_2^\star) = \arg\min_{(\hat{t}_1, \hat{t}_2) \in \mathcal{T}} \hat{t}_1 \hat{t}_2$.

4 Theoretical properties

4.1 FSDR control

Under mild conditions, the proposed false simultaneous discovery rate estimator (4) gives asymptotically conservative estimation of FSDR even when the F_{kj}^1 differ across j:

Theorem 1 *Let* $\mathrm{FSDR}(t_1, t_2)$ *be the true false simultaneous discovery rate* (1) *for the rectangular rejection region* $[0, t_1] \times [0, t_2]$. *Under model* (2) *and Condition 1 in the Appendix, if* $\pi_{11} > \pi_1\pi_2$, *then*

$$\lim_{p\to\infty} \inf_{t_1\geq\delta_1, t_2\geq\delta_2} \left\{ \widehat{\mathrm{FSDR}}(t_1, t_2) - \mathrm{FSDR}(t_1, t_2) \right\} \geq 0$$

almost surely, for each $\delta_1, \delta_2 > 0$.

The condition that $\pi_{11} > \pi_1\pi_2$ is necessary for inequality (3), and essentially states that if a given feature is significant in study 1, it is more likely to also be significant in study 2 as compared to a randomly chosen one from study 2. When the P_{kj} are interpreted as random draws from the marginal distribution F_k, this condition implies that P_{1j} and P_{2j} are positively dependent. Intuitively, identifying simultaneous signals is only sensible if the studies being analyzed are similar in some way, and positive dependence is a measure of this similarity. When this condition does not hold, $\widehat{\mathrm{FSDR}}(t_1, t_2)$ will in general be larger than 1 and will remain conservative.

As a consequence of Theorem 1, any of the optimal thresholds contained in $\mathcal{T}$ (5), including the proposed $(\hat{t}_1^\star, \hat{t}_2^\star)$ (6), offer asymptotic conservative control of FSDR. Define the pointwise limit of $\widehat{\mathrm{FSDR}}$:

$$\mathrm{FSDR}_\infty(t_1, t_2) = F_1(t_1)F_2(t_2)/G(t_1, t_2), \tag{8}$$

where G is the probability limit of $\hat{G}$ and the F_k are essentially those used in (3). These quantities are more carefully defined in (13) and (14) in the Appendix.

Theorem 2 *Let* $\mathrm{FSDR}(t_1, t_2)$ *be the true false simultaneous discovery rate* (1) *for the rejection region* $[0, t_1] \times [0, t_2]$. *Under model* (2) *and Condition 1 in the Appendix, if* $\pi_{11} > \pi_1\pi_2$ *and there exists a* (t_1', t_2') *such that* $t_1', t_2' > 0$ *and* $\mathrm{FSDR}_\infty(t_1', t_2') < \alpha$, *then for any* $(\hat{t}_1, \hat{t}_2) \in \mathcal{T}$ (5),

$$\limsup_{p\to\infty} \mathrm{FSDR}(\hat{t}_1, \hat{t}_2) \leq \alpha.$$

The proofs of Theorems 1 and 2 given in the Appendix assume that the observed data are p-values that follow model (2), but with a few minor modifications they can also be shown to hold in the setting of Section 3.6, when the null distributions of the observed test statistics are unknown.

4.2 Identifiability

It is very surprising that the proposed procedure does not require knowledge of either F_{kj}^0 or F_{kj}^1. If the test statistics are viewed from the random effects perspective, as independent draws from some two-group mixture distribution, the components of the mixture are not identifiable if neither the null nor the alternative are specified. Combined with the non-uniqueness of the optimal rectangular rejection region discussed in Section 3.4, this raises the question of whether the two-sequence simultaneous discovery problem is well-defined.

The limiting FSDR control problem is indeed identifiable under certain reasonable conditions. This is a general result that does not depend on either model (2) or model (7). Let $G(t_1, t_2)$ be any bivariate distribution function and denote its marginals by $F_k(t_k), k = 1, 2$.

Assumption 1 *For $k = 1, 2$, with $f_k(t_k) = dF_k(t_k)/dt_k$, the function*

$$t_k \mapsto \frac{\partial G(t_1, t_2)/\partial F_k(t_k)}{f_k(t_k)}$$

is strictly decreasing in t_k.

Theorem 3 *Under Assumption 1, the constrained optimization problem*

$$\underset{t_1, t_2 \in (-\infty, \infty)}{\arg\max} \left\{ G(t_1, t_2) : \frac{F_1(t_1) F_2(t_2)}{G(t_1, t_2)} \leq \alpha \right\} \tag{9}$$

has a unique solution.

When G is the probability limit of $p^{-1} \sum_j I(P_{1j} \leq t_1, P_{2j} \leq t_2)$, Theorem 3 establishes identifiability of the proposed FSDR control procedure for p-values following model (2). Similarly, defining G as the probability limit of $p^{-1} \sum_j I(-T_{1j} \leq t_1, -T_{2j} \leq t_2)$ establishes identifiability for general test statistics following model (7).

Assumption 1 has a very intuitive interpretation when $G(t_1, t_2)$ is viewed as describing the joint distribution from which test statistics (X_{1j}, X_{2j}) are randomly drawn, where X_{kj} equals either P_{kj} or $-T_{kj}$. The assumption is then equivalent to stating that $\mathrm{P}(X_{1j} \leq t_1 \mid X_{2j} = t_2)$ and $\mathrm{P}(X_{2j} \leq t_2 \mid X_{1j} = t_1)$ are strictly decreasing in t_2 and t_1, respectively. Thus Assumption 1 requires a type of positive dependence between X_{1j} and X_{2j}, which is consistent with the positive dependence condition $\pi_{11} > \pi_1 \pi_2$ necessary for Theorems 1 and 2. This particular type of dependence has been termed "positive regression dependence" or "stochastic monotonicity", and is well-studied in copula theory (Nelsen, 2013). An interesting technical point is that in this two-sequence problem, Assumption 1 plays the same role that the monotone likelihood ratio assumption (Sun and Cai, 2007; Cao et al., 2013), or the alternative distribution concavity assumption (Genovese and Wasserman, 2002; Storey, 2002; Genovese and Wasserman, 2004), play in the one-sequence false discovery rate control problem. Like those assumptions, it can be weakened, and only needs to hold for small t_1 and t_2 when the desired FSDR level α is small.

5 Simulations

5.1 Discovery using p-values

Simulations were conducted with $p = 10,000$ features. The numbers of simultaneous signals and non-null signals in each sequence were varied in different simulation settings. The P_{kj} were simulated by first generating $T_{kj} \mid I_{kj} = 0 \sim N(0, 1)$ and $T_{kj} \mid I_{kj} = 1 \sim N(\mu_{kj}, \sigma_{kj}^2)$, and then setting $P_{kj} = 2\Phi(-|T_{kj}|)$. The μ_{kj} and σ_{kj} varied across k and j and were centered around 3 and 1, respectively, and I_{kj}, μ_{kj}, and σ_{kj}^2 were fixed across all replications. The desired FSDR was set to $\alpha = 0.05$ and each simulation setting was repeated 400 times.

The proposed method was implemented with $m_1 = m_2 = 10,000$, with m_1 and m_2 defined in Section 3.5. Three alternatives procedures, described in Section 2, were also implemented; these were chosen for the availability and efficiency of their implementations. All methods were implemented at nominal FSDR levels of 0.05.

1. The Benjamini-Hochberg (BH) procedure was applied to $P_j = P_{1j} \vee P_{2j}$.

2. The method of Chung et al. (2014), which assumes that the P_{kj} are Beta-distributed under the alternative.

3. The empirical Bayes method of Heller and Yekutieli (2014).

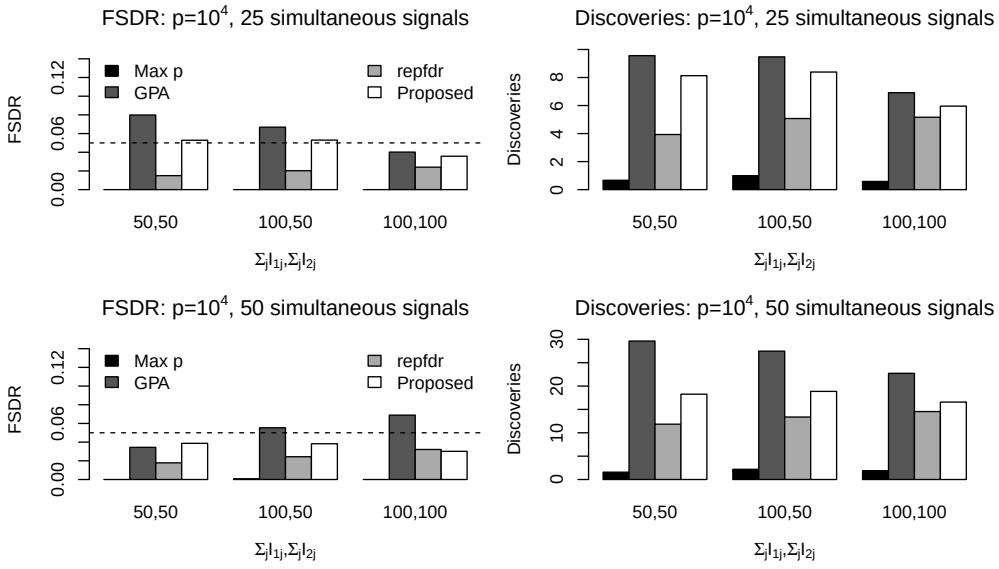

Figure 2: Empirical FSDRs and average numbers of discoveries at nominal 0.05 level, over 400 replications. Horizontal dashed lines mark the nominal 0.05 FSDR level. Max p: apply BH to $P_{1j} \vee P_{2j}$; GPA: method of Chung et al. (2014); repfdr: method of Heller and Yekutieli (2014).

Figure 2 reports the empirical FSDRs and the average numbers of discovered putative simultaneous signals. The max p method was too conservative, the method of Chung et al. (2014) was unable to maintain the nominal FSDR level, and the method of Heller and Yekutieli (2014) had low power. In contrast, our proposed procedure maintained FSDR control while still discovering many simultaneous signals.

5.2 Discovery using test statistics with unknown null distributions

Instead of p-values, test statistics were generated following $T_{kj} = \tilde{T}_{kj}^\top \tilde{T}_{kj}$, where $\tilde{T}_{kj} \in \mathbb{R}^{100}$. When $I_{kj} = 0$, $\tilde{T}_{kj} \sim N(0, \Sigma)$, where the diagonals of Σ equaled 1 and the off-diagonals equaled ρ, which was varied in different simulation settings. When $I_{kj} = 1$, $\tilde{T}_{kj}$ was generated similarly except each component of the vector had marginal mean μ_{kj} and variance σ_{kj}^2. This scheme models a typical genome-wide association study analysis, where sets of correlated genetic variants are often grouped together to calculate a single test statistic. Throughout, $p = 10,000$ features and $\sum_j I_{kj} = 100$ for $k = 1, 2$. The desired FSDR was again set to $\alpha = 0.05$ and each simulation setting was repeated 400 times.

In practice the correlation structure between the components of $\tilde{T}_{kj}$ would be unknown, leading to T_{kj} with unknown nulls. Nevertheless, the max p method and the method of Chung et al. (2014) require p-values as inputs, so pseudo-p-values were calculated assuming the T_{kj} had null distributions of χ^2_{100}, which is misspecified for $\rho \neq 0$. The method of Heller and Yekutieli (2014) requires Z-scores, which were calculated as $-\Phi^{-1}(P_{kj}/2)$ and are thus also misspecified for $\rho \neq 0$. All methods were implemented at nominal FSDR levels of 0.05.

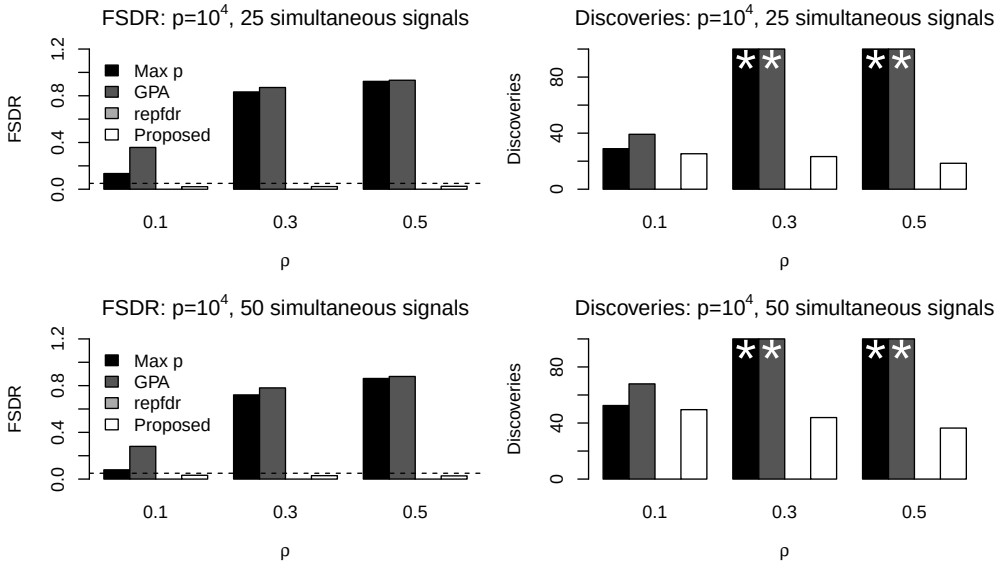

Figure 3: Empirical FSDRs and average numbers of simultaneous discoveries at nominal 0.05 level, over 400 replications. An asterisk indicates that the average number of discoveries exceeded the upper limit of the vertical axis. Horizontal dashed lines mark the nominal 0.05 FSDR level. Max p: apply BH to $P_{1j} \vee P_{2j}$; GPA: the method of Chung et al. (2014); repfdr: method of Heller and Yekutieli (2014).

Figure 3 reports the empirical FSDRs and the average numbers of discoveries. Existing methods were heavily dependent on correct specification of the null of T_{kj} and performed very poorly; the method of Heller and Yekutieli (2014) encountered numerical errors. Our procedure was the only one that controlled the FSDR while having enough power to discover nearly all of the simultaneous signals.

6 Application to pleiotropy in psychiatric disorders

There is a great deal of evidence for the role of pleiotropy, the phenomenon where genetic variants affect more than one phenotype, in human disease (Sivakumaran et al., 2011). The Cross-Disorder Group of the Psychiatric Genomics Consortium was established to study the genetic relationship between five psychiatric disorders, namely attention-deficit disorder, autism-spectrum disorders, bipolar disorder, major depressive disorder, and schizophrenia (Cross-Disorder Group of the Psychiatric Genomics Consortium et al., 2013a,b). The consortium examined 33,332 cases and 27,888 controls across these five dis-

orders, and summary p-values for 1,219,805 genotyped and imputed single-nucleotide polymorphisms are available for all five conditions from the Psychiatric Genomics Consortium website at `http://www.med.unc.edu/pgc/downloads`.

The method proposed in this paper was applied to each pair of disorders in order to discover genetic variants simultaneously associated with both conditions in the pair. The other existing methods described in the simulations in Section 5 were also applied. The nominal false simultaneous discovery rate was set to 0.05 throughout. This pairwise analysis provides a fine-resolution picture of the nature of pleiotropy in psychiatric disorders. These results can potentially lead to improvements in classification, diagnosis, and treatment.

Table 2: Number of discovered pleiotropic variants; ADD = attention-deficit disorder, AUT = autism-spectrum disorders, BIP = bipolar disorder, MDD = major depressive disorder, SCZ = schizophrenia. Max p: apply BH to $P_{1j} \vee P_{2j}$; GPA: method of Chung et al. (2014); repfdr: method of Heller and Yekutieli (2014). Dashes indicate numerical errors encountered by the software.

Disorder 1	Disorder 2	Max p	GPA	repfdr	Proposed
ADD	AUT	0	0	–	0
ADD	BIP	0	0	–	0
ADD	MDD	0	0	–	0
ADD	SCZ	0	0	–	0
AUT	BIP	0	0	0	0
AUT	MDD	0	0	–	5
AUT	SCZ	0	0	0	0
BIP	MDD	0	0	–	72
BIP	SCZ	0	359	15	128
MDD	SCZ	0	0	–	15

Table 2 reports the number of potentially pleiotropic variants discovered by each method. Three out of the four methods identified numerous pleiotropic variants associated with both bipolar disorder and schizophrenia, which is consistent with previous findings (Andreassen et al., 2013a). The proposed method was the only one to discover pleiotropic variants that affect both bipolar disorder and major depression, as well as both major depression and schizophrenia, suggesting a genetic connection between these three disorders. These findings are likely not false positives. First, Huang et al. (2010) obtained similar results. In addition, children with autism have been found to suffer from comorbid major depression (Leyfer et al., 2006). Finally, the five single nucleotide polymorphisms identified as simultaneously associated with autism and major depression lie in the gene SEMA3A, which is involved in axonal guidance and known to be involved in psychiatric diseases such as schizophrenia (Eastwood et al., 2003).

To assess the robustness of these findings to the correlation structure of the genetic variants, the p-values were pruned by removing those corresponding to highly correlated polymorphisms. Specifically, the `--indep-pairwise` function of the genome-wide association software Plink (Purcell et al., 2007) was used to identify variants in approximate linkage equilibrium. The Psychiatric Genomics Consortium studied only subjects with European

ancestry, so linkage disequilibrium between variants was estimated using the HapMap CEU samples, and r^2 thresholds of either 0.2 or 0.5 were used. After pruning, the proposed method was the only method that made any discoveries. Pleiotropic variants affecting bipolar disorder and schizophrenia, as well as major depression and schizophrenia, were still discovered for both r^2 thresholds, and variants affecting bipolar disorder and depression were still discovered for $r^2 = 0.5$. The major findings of the proposed method were thus fairly consistent across different sets of pruned variants.

7 Discussion

The proposed procedure has so far been stated assuming independence of test statistics across features, but this can be weakened as long as the strong law of large numbers still holds for the test statistics (Storey et al., 2004). For example, a result of Gaposhkin (1995) implies that under certain conditions, Theorems 1 and 2 will still hold for blockwise-independent test statistics. Furthermore, it is likely that the procedure is robust to stronger forms of dependence. In the single-sequence multiple testing problem with dependent test statistics, it is possible to control the false discovery proportion, which is the realized ratio of false discoveries to total discoveries (Genovese and Wasserman, 2004; Efron, 2007, 2010a; Fan et al., 2012). The key is to estimate an empirical null distribution for the test statistics rather than using the theoretical null (Schwartzman, 2012), because under dependence the two can be very different. But in the proposed procedure, $\widehat{\text{FSDR}}$ (4) already uses empirical distribution estimates, so it may also be able to control the false simultaneous discovery proportion. This will be explored in future work.

Simultaneous signal identification is distinct from meta-analysis. The latter seeks to estimate and test some overall measure of association across studies. For example, a genetic variant that is significant in some studies but not others may still have a large average cross-study effect. It would be identified by a meta-analysis, but is not a simultaneous signal. The FSDR control problem is more closely related to replicability analysis (Bogomolov and Heller, 2013; Heller et al., 2014; Heller and Yekutieli, 2014), which, given a set of significant findings from a primary study, seeks to control the false discovery rate for identifying replicated findings in follow-up study. The setting studied in this paper makes no distinction between a primary and follow-up study.

An interesting feature of the FSDR control problem is that it can sometimes be easier to identify simultaneous signals than to separately identify non-null signals in each sequence of test statistics. Intuitively, this occurs when the non-null associations in one of the sequences are very strong. Given two sequences of p-values $P_{kj}, k = 1, 2$, suppose it is known exactly which of the P_{2j} are non-null. Then to find simultaneous signals, only the P_{1j} corresponding to the non-null P_{2j} need to be considered. This greatly reduces the effective number of test pairs, which allows for larger rejection region for the P_{1j} than would be obtained if a standard false discovery rate control procedure were used on the P_{1j} alone. In other words, simultaneous discovery can lead to more powerful single-sequence discovery.

8 Appendix

8.1 Preliminaries

The proofs assume that the observed data are p-values P_{1j} and P_{2j} following model (2), but they are easily modified to accommodate any continuous test statistics following model (7). Using the notation in Table 1, define the following functions and technical condition:

$$\begin{aligned} V_{ab}(t_1, t_2) &= \sum_{j: I_{1j}=a, I_{2j}=b} I(P_{1j} \le t_1, P_{2j} \le t_2), \quad a, b = 0, 1, \\ R(t_1, t_2) &= \sum_i I(P_{i1} \le t_1, P_{i2} \le t_2). \end{aligned} \tag{10}$$

Condition 1 *As $p \to \infty$, the distribution functions F_{kj}^0 and F_{kj}^1 defined in model (2) satisfy*

$$\begin{aligned} &\frac{1}{p_{00}} \sum_{j: I_{1j}=0, I_{2j}=0} F_{kj}^0(t) \to F_k^0(t), \\ &\frac{1}{p_{10}} \sum_{j: I_{1j}=1, I_{2j}=0} F_{1j}^1(t) \to F_1^1(t), \quad \frac{1}{p_{10}} \sum_{j: I_{1j}=1, I_{2j}=0} F_{2j}^0(t) \to F_2^0(t), \\ &\frac{1}{p_{01}} \sum_{j: I_{1j}=0, I_{2j}=1} F_{1j}^0(t) \to F_1^0(t), \quad \frac{1}{p_{01}} \sum_{j: I_{1j}=0, I_{2j}=1} F_{2j}^1(t) \to F_2^1(t), \\ &\frac{1}{p_{11}} \sum_{i: I_{1j}=1, I_{2j}=1} F_{kj}^1(t) \to F_k^1(t), \quad k = 1, 2, \end{aligned}$$

pointwise in $t \in [0, 1]$ for some functions F_k^0 and F_k^1, $k = 1, 2$. In addition, $p_{ab}/p \to \pi_{ab}$ for $a, b = 0, 1$.

Condition 1 is clearly satisfied when for each k, the F_{kj}^0 and F_{kj}^1 do not vary across j. When this is not true, F_k^0 and F_k^1 can be thought of as mixture distributions. This is similar to the two-group random effects model for a sequence of hypothesis tests common in the multiple testing literature (Sun and Cai, 2007; Efron, 2010b).

8.2 Proof of Theorem 1

For a given rectangular rejection region $[0, t_1] \times [0, t_2]$, the true FSDR (1) is

$$\mathrm{FSDR}(t_1, t_2) = E\left\{ \frac{V_{00}(t_1, t_2) + V_{10}(t_1, t_2) + V_{01}(t_1, t_2)}{1 \vee R(t_1, t_2)} \right\}.$$

Define $\hat{F}_k(t) = p^{-1} \sum_j I(P_{kj} \le t), k = 1, 2$, so that

$$\widehat{\mathrm{FSDR}}(t_1, t_2) = \frac{p\hat{F}_1(t_1)\hat{F}_2(t_2)}{1 \vee R(t_1, t_2)}.$$

It will first be shown that for $\delta_1, \delta_2 > 0$,

$$\lim_{p\to\infty} \inf_{t_1\geq\delta_1, t_2\geq\delta_2} \left\{ \widehat{\mathrm{FSDR}}(t_1,t_2) - \frac{V_{00}(t_1,t_2)+V_{10}(t_1,t_2)+V_{01}(t_1,t_2)}{1\vee R(t_1,t_2)} \right\} \geq 0 \tag{11}$$

almost surely. Next it will be shown that

$$\sup_{t_1\geq\delta_1, t_2\geq\delta_2} \left| \frac{V_{00}(t_1,t_2)+V_{10}(t_1,t_2)+V_{01}(t_1,t_2)}{1\vee R(t_1,t_2)} - \mathrm{FSDR}(t_1,t_2) \right| \xrightarrow{a.s.} 0, \tag{12}$$

which will complete the proof.

Inequality (11) is true if and only if

$$\lim_{p\to\infty} \inf_{t_1\geq\delta_1, t_2\geq\delta_2} \{ p\hat{F}_1(t_1)\hat{F}_2(t_2) - V_{00}(t_1,t_2) - V_{10}(t_1,t_2) - V_{01}(t_1,t_2) \} \geq 0.$$

Using a version of the strong law of large numbers and the bivariate Glivenko-Cantelli theorem (Pollard, 1984, Ch. II), it can be shown that under model (2) and Condition 1 that

$$\sup_{t_1,t_2\in[0,1]} |p^{-1}V_{ab}(t_1,t_2) - \pi_{ab}F_1^a(t_1)F_2^b(t_2)| \xrightarrow{a.s.} 0$$

for $a, b = 0, 1$, π_k, F_k^0, and F_k^1 as defined in Condition 1. Defining

$$\pi_k = p^{-1}\sum_j I_{kj}, \quad F_k(t) = (1-\pi_k)F_k^0(t) + \pi_k F_k^1(t), \quad k = 1, 2, \tag{13}$$

this implies that

$$\sup_{t_1\in[0,1]} |\hat{F}_k(t) - F_k(t)| \xrightarrow{a.s.} 0, \quad k = 1, 2.$$

In addition,

$$\begin{aligned}
&\sup_{t_1,t_2\in[0,1]} |\hat{F}_1(t_1)\hat{F}_2(t_2) - F_1(t_1)F_2(t_2)| \\
\leq &\sup_{t_1,t_2\in[0,1]} |\hat{F}_1(t_1)\hat{F}_2(t_2) - \hat{F}_1(t_1)F_2(t_2)| + \sup_{t_1,t_2\in[0,1]} |\hat{F}_1(t_1)F_2(t_2) - F_1(t_1)F_2(t_2)| \\
\leq &\sup_{t_2\in[0,1]} |\hat{F}_2(t_2) - F_2(t_2)| + \sup_{t_1\in[0,1]} |\hat{F}_1(t_1) - F_1(t_1)| \xrightarrow{a.s.} 0.
\end{aligned}$$

Together with (3), these results imply (11).

To show (12), first define

$$\begin{aligned}
G_0(t_1,t_2) &= \pi_{00}F_1^0(t_1)F_2^0(t_2) + \pi_{10}F_1^1(t_1)F_2^0(t_2) + \pi_{01}F_1^0(t_1)F_2^1(t_2), \\
G(t_1,t_2) &= G_0(t_1,t_2) + \pi_{11}F_1^1(t_1)F_2^1(t_2).
\end{aligned} \tag{14}$$

By the triangle inequality

$$
\begin{aligned}
0 \leq & \sup_{t_1 \geq \delta_1, t_2 \geq \delta_2} \left| \frac{V_{00}(t_1,t_2) + V_{10}(t_1,t_2) + V_{01}(t_1,t_2)}{1 \vee R(t_1,t_2)} - \frac{G_0(t_1,t_2)}{G(t_1,t_2)} \right| \\
\leq & \sup_{t_1 \geq \delta_1, t_2 \geq \delta_2} \left| \frac{V_{00}(t_1,t_2) + V_{10}(t_1,t_2) + V_{01}(t_1,t_2)}{1 \vee R(t_1,t_2)} - \frac{pG_0(t_1,t_2)}{1 \vee R(t_1,t_2)} \right| + \\
& \sup_{t_1 \geq \delta_1, t_2 \geq \delta_2} \left| \frac{pG_0(t_1,t_2)}{1 \vee R(t_1,t_2)} - \frac{G_0(t_1,t_2)}{G(t_1,t_2)} \right| \\
\leq & \frac{p}{1 \vee R(\delta_1,\delta_2)} \sup_{t_1 \geq \delta_1, t_2 \geq \delta_2} \left| \frac{V_{00}(t_1,t_2) + V_{10}(t_1,t_2) + V_{01}(t_1,t_2)}{p} - G_0(t_1,t_2) \right| + \\
& \frac{pG_0(1,1)}{\{1 \vee R(\delta_1,\delta_2)\} G(\delta_1,\delta_2)} \sup_{t_1 \geq \delta_1, t_2 \geq \delta_2} |G(t_1,t_2) + p^{-1}\{1 \vee R(t_1,t_2)\}| \\
\xrightarrow{a.s.} & 0.
\end{aligned}
$$

Therefore

$$
\begin{aligned}
0 = & \mathrm{E} \lim_{m \to \infty} \sup_{t_1 \geq \delta_1, t_2 \geq \delta_2} \left| \frac{V_{00}(t_1,t_2) + V_{0A}(t_1,t_2) + V_{A0}(t_1,t_2)}{1 \vee R(t_1,t_2)} - \frac{G_0(t_1,t_2)}{G(t_1,t_2)} \right| \\
= & \lim_{m \to \infty} \sup_{t_1 \geq \delta_1, t_2 \geq \delta_2} \mathrm{E} \left| \frac{V_{00}(t_1,t_2) + V_{0A}(t_1,t_2) + V_{A0}(t_1,t_2)}{1 \vee R(t_1,t_2)} - \frac{G_0(t_1,t_2)}{G(t_1,t_2)} \right| \\
\geq & \lim_{m \to \infty} \sup_{t_1 \geq \delta_1, t_2 \geq \delta_2} \left| \mathrm{FSDR}(t_1,t_2) - \frac{G_0(t_1,t_2)}{G(t_1,t_2)} \right| \geq 0,
\end{aligned}
$$

where the second equality follows from the dominated convergence theorem. Together these results imply (12).

8.3 Proof of Theorem 2

The theorem is trivially true when $(\hat{t}_1, \hat{t}_2) = (0,0)$. Otherwise, suppose there exist $\delta_1, \delta_2 > 0$ such that $\hat{t}_1 \geq \delta_1$ and $\hat{t}_2 \geq \delta_2$ for any $(\hat{t}_1, \hat{t}_2) \in \mathcal{T}$. This would imply that

$$
\begin{aligned}
\lim_{p \to 0} \{\alpha - \mathrm{FSDR}(\hat{t}_1, \hat{t}_2)\} &\geq \lim_{p \to 0} \{\widehat{\mathrm{FSDR}}(\hat{t}_1, \hat{t}_2) - \mathrm{FSDR}(\hat{t}_1, \hat{t}_2)\} \\
&\geq \lim_{p \to 0} \inf_{t_1 \geq \delta_1, t_2 \geq \delta_2} \{\widehat{\mathrm{FSDR}}(t_1,t_2) - \mathrm{FSDR}(t_1,t_2)\} \geq 0
\end{aligned}
$$

where the last inequality follows from Theorem 1. The desired result follows.

To construct δ_1 and δ_2, by assumption there exists an $\epsilon > 0$ such that $\mathrm{FSDR}_\infty(t_1', t_2') = \alpha - \epsilon$, with FSDR defined in (8). Using a version of the strong law of large numbers and Slutsky's theorem, it can be shown that for p sufficiently large,

$$
|\widehat{\mathrm{FSDR}}(t_1', t_2') - \mathrm{FSDR}(t_1', t_2')| \leq \epsilon/2
$$

with probability 1, which implies that $\widehat{\mathrm{FSDR}}(t_1', t_2') \leq \alpha - \epsilon/2$. In other words, (t_1', t_2') is a feasible solution of the optimization problem (5).

Next define $\hat{y}(t_1)$ and $y'(t_1)$ to be level curves of the empirical bivariate distribution function $\hat{G}$ such that for all $t_1 \in [0,1]$,

$$\hat{G}\{t_1, \hat{y}(t_1)\} = \hat{G}(\hat{t}_1, \hat{t}_2), \quad \hat{G}\{t_1, y'(t_1)\} = \hat{G}(t_1', t_2').$$

Since (t_1', t_2') is a feasible solution of the optimization problem, $\hat{G}(\hat{t}_1, \hat{t}_2) \geq \hat{G}(t_1', t_2')$ for any $(\hat{t}_1, \hat{t}_2) \in \mathcal{T}$. Since $\hat{G}(t_1, t_2)$ is a nondecreasing function of t_2, $\hat{y}(t_1) \geq y'(t_1)$ for all $t_1 \in [0,1]$. Define $\delta_2 = \inf_{t_1} y'(t_1)$. Then

$$\hat{G}(\hat{t}_1, \hat{t}_2) = \hat{G}\{\hat{t}_1, y'(\hat{t}_1)\} \geq \hat{G}(\hat{t}_1, \delta_2),$$

which implies that $\hat{t}_2 \geq \delta_2$. A similar line of reasoning can be employed to construct δ_1.

It remains to show that for m sufficiently large, $\delta_1, \delta_2 > 0$ with probability 1. This will be shown for δ_2; the proof for δ_1 is similar. Since t_1' and t_2' are both positive, $G(t_1', t_2') = \epsilon > 0$, where $G(t_1, t_2)$ is defined in (14). By a version of the strong law of large numbers and Condition 1, it can be shown that for p sufficiently large $\hat{G}(t_1', t_2') > \epsilon/2$ with probability 1. Therefore for all $t_1 \in [0,1]$, $\hat{G}\{t_1, y'(t_1)\} = \hat{G}(t_1', t_2') > \epsilon/2$. If δ_2 were equal to zero, then $\inf_{t_1} \hat{G}\{t_1, y'(t_1)\} = 0$, which leads to a contradiction.

8.4 Proof of Theorem 3

By Sklar's theorem, there exists a unique copula $C(u,v)$ such that $C(u,v) = G\{F_1^{-1}(u), F_2^{-1}(v)\}$ for $u, v \in [0,1]$. Define $D(u,v) = uv/C(u,v)$. Then problem (9) is equivalent to

$$\arg\max_{u,v \in [0,1]} C(u,v) \text{ such that } D(u,v) \leq \alpha \tag{15}$$

and letting $t_1 = F_1^{-1}(u)$ and $t_2 = F_2^{-1}(v)$. It remains to prove that (15) has a unique solution.

First it will be shown that the solution must lie on the level curve $L_\alpha(u)$ defined by $D\{u, L_\alpha(u)\} = \alpha$. By Assumption (1), $\partial C(u,v)/\partial u$ is strictly decreasing in u, so

$$C(u,v) = \int_0^u \left.\frac{\partial C}{\partial u}\right|_{(x,v)} dx > \left.\frac{\partial C}{\partial u}\right|_{(u,v)} \int_0^u dx = u \left.\frac{\partial C}{\partial u}\right|_{(u,v)}$$

for any (u,v). Therefore

$$\frac{\partial D}{\partial u} = \frac{v}{C(u,v)^2}\left\{C(u,v) - u\frac{\partial C}{\partial u}\right\} > 0,$$

and by similar reasoning $\partial D/\partial v > 0$ for any (u,v). In other words, $D(u,v)$ is strictly increasing in u and v. Thus for any (u',v') such that $D(u',v') < \alpha$, there exists a (u,v) with either $u > u'$ or $v > v'$ such that $D(u',v') = \alpha$. Since in addition $C(u,v) > C(u',v')$, (u',v') cannot be a solution to (15).

Next it will be shown that the function $u \mapsto C\{u, L_\alpha(u)\}$ is strictly concave. Since

$$\frac{d^2}{du^2} C\{u, L_\alpha(u)\} = \alpha^{-1}\frac{d^2}{du^2}\{uL_\alpha(u)\} = \alpha^{-1}\{uL_\alpha''(u) + 2L_\alpha'(u)\},$$

it suffices to show that $L''_\alpha(u) < -2L'_\alpha(u)/u$. By definition, the vector tangent to $L_\alpha(u)$ at u is perpendicular to the gradient of $D(u,v)$ at u. It was established above that $\partial D/\partial u$ and $\partial D/\partial u$ are positives, so

$$\frac{dL_\alpha(u)}{du} = -\frac{\partial D/\partial u|_{\{u,f(u)\}}}{\partial D/\partial v|_{\{u,f(u)\}}} = -\frac{L_\alpha(u) - \alpha h_1(u)}{u - \alpha h_2(u)},$$

where

$$h_1(u) = \left.\frac{\partial C}{\partial u}\right|_{\{u,f(u)\}}, \quad h_2(u) = \left.\frac{\partial C}{\partial v}\right|_{\{u,f(u)\}}.$$

Then

$$\begin{aligned}
\frac{d^2L_\alpha(u)}{du^2} &= -\frac{\{u - \alpha h_2(u)\}\{L'_\alpha(u) - \alpha h'_1(u)\} - \{L_\alpha(u) - \alpha h_1(u)\}\{1 - \alpha h'_2(u)\}}{\{u - \alpha h_2(u)\}^2} \\
&= -\frac{L'_\alpha(u) - \alpha h'_1(u) + L'_\alpha(u)\{1 - \alpha h'_2(u)\}}{u - \alpha h_2(u)} \\
&= -\frac{2L'_\alpha(u)}{u - \alpha h_2(u)} + \alpha\frac{h'_1(u) + L'_\alpha(u)h'_2(u)}{u - \alpha h_2(u)}.
\end{aligned}$$

It can be shown that the second term on the right-hand side above is negative. Since $u - \alpha h_2(u) = \partial D/\partial v > 0$, it remains to show that $h'_1(u) + L'_\alpha(u)h'_2(u) < 0$. First, L_α must be decreasing in u. For any $u' > u$, because $D(u,v)$ is increasing in u and v, $D\{u, L_\alpha(u)\} < D\{u', L_\alpha(u)\}$. For L_α to be a level curve of D, it must be that $L_\alpha(u') < L_\alpha(u)$, so $L'_\alpha(u) < 0$. Next,

$$h_1(u) = \left.\frac{\partial C}{\partial u}\right|_{u,L_\alpha(u)} > \left.\frac{\partial C}{\partial u}\right|_{u',L_\alpha(u)} \geq \left.\frac{\partial C}{\partial u}\right|_{u',L_\alpha(u')} = h_1(u'),$$

where the first inequality is due to Assumption 1 and the second is because the function $v \mapsto \partial C(u,v)/\partial u$ is nondecreasing (Nelsen, 2013). Thus $h_1(u)$ is decreasing as well and $h'_1(u) < 0$. Similarly,

$$h_2(u) = \left.\frac{\partial C}{\partial v}\right|_{u,L_\alpha(u)} < \left.\frac{\partial C}{\partial v}\right|_{u,L_\alpha(u')} \leq \left.\frac{\partial C}{\partial v}\right|_{u',L_\alpha(u')} = h_2(u'),$$

so $h_2(u)$ is increasing and $h'_2(u) > 0$. Combining these results gives $h'_1(u) + L'_\alpha(u)h'_2(u) < 0$.

Finally, $\{u - \alpha h_2(u)\}^{-1} \geq u^{-1}$ because $u - \alpha h_2(u) > 0$ and the partial derivatives of a copula are always positive (Nelsen, 2013), implying that

$$L''_\alpha(u) < -\frac{2L'_\alpha(u)}{u} + \alpha\frac{h'_1(u) + L'_\alpha(u)h'_2(u)}{u - \alpha h_2(u)} < -\frac{2L'_\alpha(u)}{u}.$$

Therefore $d^2C\{u, L_\alpha(u)\}/du^2 < 0$ and (15) has a unique maximum.

References

O. A. Andreassen, S. Djurovic, W. K. Thompson, A. J. Schork, K. S. Kendler, M. C. ODonovan, D. Rujescu, T. Werge, M. van de Bunt, A. P. Morris, et al. Improved detection of common variants associated with schizophrenia by leveraging pleiotropy with

cardiovascular-disease risk factors. *The American Journal of Human Genetics*, 92(2): 197–209, 2013a.

O. A. Andreassen, W. K. Thompson, A. J. Schork, S. Ripke, M. Mattingsdal, J. R. Kelsoe, K. S. Kendler, M. C. O'Donovan, D. Rujescu, T. Werge, et al. Improved detection of common variants associated with schizophrenia and bipolar disorder using pleiotropy-informed conditional false discovery rate. *PLoS genetics*, 9(4):e1003455, 2013b.

Y. Benjamini and Y. Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 289–300, 1995.

M. Bogomolov and R. Heller. Discovering findings that replicate from a primary study of high dimension to a follow-up study. *Journal of the American Statistical Association*, 108 (504):1480–1492, 2013.

H. Cao, W. Sun, and M. R. Kosorok. The optimal power puzzle: scrutiny of the monotone likelihood ratio assumption in multiple testing. *Biometrika*, page ast001, 2013.

Z. Chi. False discovery rate control with multivariate p-values. *Electronic Journal of Statistics*, 2:368–411, 2008.

D. Chung, C. Yang, C. Li, J. Gelernter, and H. Zhao. Gpa: A statistical approach to prioritizing gwas results by integrating pleiotropy and annotation. *PLoS genetics*, 10(11): e1004787, 2014.

Cross-Disorder Group of the Psychiatric Genomics Consortium et al. Genetic relationship between five psychiatric disorders estimated from genome-wide snps. *Nature genetics*, 2013a.

Cross-Disorder Group of the Psychiatric Genomics Consortium et al. Identification of risk loci with shared effects on five major psychiatric disorders: a genome-wide analysis. *Lancet*, 381(9875):1371, 2013b.

D. Donoho. 50 years of data science. Technical report, Stanford University, September 2015.

D. Donoho and J. Jin. Higher criticism for detecting sparse heterogeneous mixtures. *The Annals of Statistics*, 32(3):962–994, 2004.

L. Du and C. Zhang. Single-index modulated multiple testing. *The Annals of Statistics*, 42 (4):30–79, 2014.

S. Eastwood, A. Law, I. Everall, and P. Harrison. The axonal chemorepellant semaphorin 3a is increased in the cerebellum in schizophrenia and may contribute to its synaptic pathology. *Molecular psychiatry*, 8(2):148–155, 2003.

B. Efron. Correlation and large-scale simultaneous significance testing. *Journal of the American Statistical Association*, 102(477), 2007.

B. Efron. Correlated z-values and the accuracy of large-scale statistical estimates. *Journal of the American Statistical Association*, 105(491):1042–1055, 2010a.

B. Efron. *Large-scale inference: empirical Bayes methods for estimation, testing, and prediction*, volume 1. Cambridge University Press, 2010b.

J. Fan, X. Han, and W. Gu. Estimating false discovery proportion under arbitrary covariance dependence. *Journal of the American Statistical Association*, 107(499):1019–1035, 2012.

V. F. Gaposhkin. On the strong law of large numbers for blockwise independent and blockwise orthogonal random variables. *Theory of Probability & Its Applications*, 39(4):677–684, 1995.

C. Genovese and L. Wasserman. Operating characteristics and extensions of the false discovery rate procedure. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(3):499–517, 2002.

C. Genovese and L. Wasserman. A stochastic process approach to false discovery control. *Annals of Statistics*, pages 1035–1061, 2004.

R. Heller and D. Yekutieli. Replicability analysis for genome-wide association studies. *The Annals of Applied Statistics*, 8(1):481–498, 2014.

R. Heller, M. Bogomolov, and Y. Benjamini. Deciding whether follow-up studies have replicated findings in a preliminary large-scale omics study. *Proceedings of the National Academy of Sciences*, 111(46):16262–16267, 2014.

J. Huang, R. H. Perlis, P. H. Lee, A. J. Rush, M. Fava, G. S. Sachs, J. Lieberman, S. P. Hamilton, P. Sullivan, P. Sklar, et al. Cross-disorder genomewide analysis of schizophrenia, bipolar disorder, and depression. *American Journal of Psychiatry*, 167(10):1254–1263, 2010.

J. P. Ioannidis, E. E. Ntzani, T. A. Trikalinos, and D. G. Contopoulos-Ioannidis. Replication validity of genetic association studies. *Nature genetics*, 29(3):306–309, 2001.

O. T. Leyfer, S. E. Folstein, S. Bacalman, N. O. Davis, E. Dinh, J. Morgan, H. Tager-Flusberg, and J. E. Lainhart. Comorbid psychiatric disorders in children with autism: Interview development and rates of disorders. *Journal of autism and developmental disorders*, 36(7):849–861, 2006.

Y. R. Li, J. Li, S. D. Zhao, J. P. Bradfield, F. D. Mentch, S. M. Maggadottir, C. Hou, D. J. Abrams, D. Chang, F. Gao, D. Guo, Z. Wei, J. J. Connoly, C. C., M. Bakay, C. Glessner, J.and Kao, K. A. Thomas, H. Qiu, R. Chiavacci, C. Kim, F. Wang, M. D. Snyder, J.and Richie, B. Flatø, Ø. Førre, L. Denson, S. D. Thompson, M. L. Becker, S. L. Guthery, A. Latiano, E. Perez, E. Resnick, R. D. Russell, D. Wilson, M. S. Silverberg, V. Annese, B. A. Lie, M. Punaro, M. C. Dubinsky, D. S. Monos, C. Strisciuglio, A. Staiano, E. Miele, S. Kugathasan, J. A. Ellis, J. Munro, K. Sullivan, C. Wise, H. Chapel, C. Cunningham-Rundles, S. F. A. Grant, J. Orange, P. M. A. Sleiman, E. Behrens, A. Griffiths, J. Satsangi,

T. Finkel, A. Keinan, E. T. Luning Prak, C. Polychronakos, B. Baldassano, H. Li, B. J. Keating, and H. Hakonarson. Meta-analysis of shared genetic architecture in pediatric autoimmune disease. *Nature Medicine*, 2015a. in press.

Y. R. Li, S. D. Zhao, M. Mohebnasab, J. Li, J. Bradfield, L. Steel, D. Abrams, J. Kobie, F. Mentch, J. Glessner, Y. Guo, Z. Wei, C. Cardinale, M. Bakay, J. Connoly, D. Li, S. M. Maggadottir, K. A. Thomas, H. Qiu, R. Chiavacci, C. Kim, F. Wang, J. Snyder, B. Flatø, Ø. Førre, L. Denson, S. D. Thompson, M. Becker, S. L. Guthery, A. Latiano, E. Perez, E. Resnick, C. Strisciuglio, A. Staiano, E. Miele, M. Silverberg, B. A. Lie, M. Punaro, R. Russell, D. Wilson, M. C. Dubinsky, D. S. Monos, V. Annese, J. Munro, C. Wise, H. Chapel, C. Cunningham-Rundles, J. Orange, E. M. Behrens, K. Sullivan, S. Kugathasan, A. Griffiths, J. Satsangi, S. Grant, P. Sleiman, T. Finkel, C. Polychronakos, R. N. Baldassano, E. Luning Prak, J. Ellis, H. Li, B. J. Keating, and H. Hakonarson. Genetic sharing and heritability of pediatric age of onset autoimmune diseases. *Nature Communications*, 2015b. in press.

R. B. Nelsen. *An introduction to copulas*, volume 139. Springer Science & Business Media, 2013.

M. A. Newton, A. Noueiry, D. Sarkar, and P. Ahlquist. Detecting differential gene expression with a semiparametric hierarchical mixture method. *Biostatistics*, 5(2):155–176, 2004.

D. Phillips and D. Ghosh. Testing the disjunction hypothesis using voronoi diagrams with applications to genetics. *The Annals of Applied Statistics*, 8(2):801–823, 2014.

D. Pollard. *Convergence of stochastic processes*. Springer Science & Business Media, 1984.

S. Purcell, B. Neale, K. Todd-Brown, L. Thomas, M. A. Ferreira, D. Bender, J. Maller, P. Sklar, P. I. De Bakker, M. J. Daly, et al. Plink: a tool set for whole-genome association and population-based linkage analyses. *The American Journal of Human Genetics*, 81(3): 559–575, 2007.

T. Rankinen, M. A. Sarzynski, S. Ghosh, and C. Bouchard. Are there genetic paths common to obesity, cardiovascular disease outcomes, and cardiovascular risk factors? *Circulation Research*, 116(5):909–922, 2015.

A. Schwartzman. Comment: Fdp vs fdr and the effect of conditioning. *Journal of the American Statistical Association*, 107(499):1039–1041, 2012.

S. Sivakumaran, F. Agakov, E. Theodoratou, J. G. Prendergast, L. Zgaga, T. Manolio, I. Rudan, P. McKeigue, J. F. Wilson, and H. Campbell. Abundant pleiotropy in human complex diseases and traits. *The American Journal of Human Genetics*, 89(5):607–618, 2011.

J. D. Storey. A direct approach to false discovery rates. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(3):479–498, 2002.

J. D. Storey and R. Tibshirani. Statistical significance for genomewide studies. *Proceedings of the National Academy of Sciences*, 100(16):9440–9445, 2003.

J. D. Storey, J. E. Taylor, and D. Siegmund. Strong control, conservative point estimation and simultaneous conservative consistency of false discovery rates: a unified approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66(1):187–205, 2004.

W. Sun and T. T. Cai. Oracle and adaptive compound decision rules for false discovery rate control. *Journal of the American Statistical Association*, 102(479):901–912, 2007.

P. M. Visscher and G. W. Montgomery. Genome-wide association studies and human disease: from trickle to flood. *Journal of the American Medical Association*, 302(18):2028–2029, 2009.